MASSACHUSETTS INSTITUTE OF TECHNOLOGY

A.I. LABORATORY

Artificial Intelligence
Memo No. 348

December, 1975
LOGO MEMO 21

TURTLE ESCAPES THE PLANE:
SOME ADVANCED TURTLE GEOMETRY

by

Andy diSessa

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the Laboratory's education research is provided in part by the National Science Foundation under Grant EC-40708X.

TURTLE ESCAPES THE PLANE: SOME ADVANCED TURTLE GEOMETRY

Andrea A. diSessa

December 12, 1975

ABSTRACT

Since the LOGO Turtle took his first step he has been mathematically confined to running around on flat surfaces. Fortunately the physically intuitive, procedurally oriented nature of the Turtle which makes him a powerful explorer in the plane is equally, if not more apparent when he is liberated to tread curved surfaces. This paper is aimed roughly at the High School level. Yet because it is built on intuition and physical action rather than formalism, it can reach such "graduate school" mathematical ideas as geodesics, Gaussian Curvature, and topological invariants as expressed in the Gauss-Bonnet Theorem.

I. TRIANGLES

This paper is an exploration into the dark and dangerous continent of mathematics wherein we shall trek from the almost civilized land of geometry, through the forbidden grounds of differential geometry and thence to topology where many a soul has perished on the great, barren and infinitely extensible rubber sheets. I think, however, I have chosen for you a path which will show you some of the great sights without undue physical danger. In fact I would be greatly disappointed if you do not return safely and with a great store of souvenirs to entice you to return on your own and explore for yourself.

Let me begin with a humble triangle; any one will do. Everyone knows a triangle has 180° worth of angles in its three vertices. What a wonderful thing that any triangle, no matter how big or small or how it is shaped, has exactly 180°! How do you know that is true? I don't think it's obvious. After all, the angles are in different places. Let me show you my favorite way to sum the angles in a triangle.

I have a turtle who can be an excellent guide to many places in mathematics. Yet he can only do two things, walk in a straight line or turn through any angle. Luckily he is smart enough to tell the measure of any angle when he turns through it. So I can send him out to measure the angles in lots of triangles, one at a time and find out what the sum of them is.

After watching Turtle measuring triangles for a while I noticed he always does the same thing. He starts at vertex 1 of a triangle and aims toward vertex 3. Then he turns toward vertex 2 and therefore measures the angle in vertex 1.

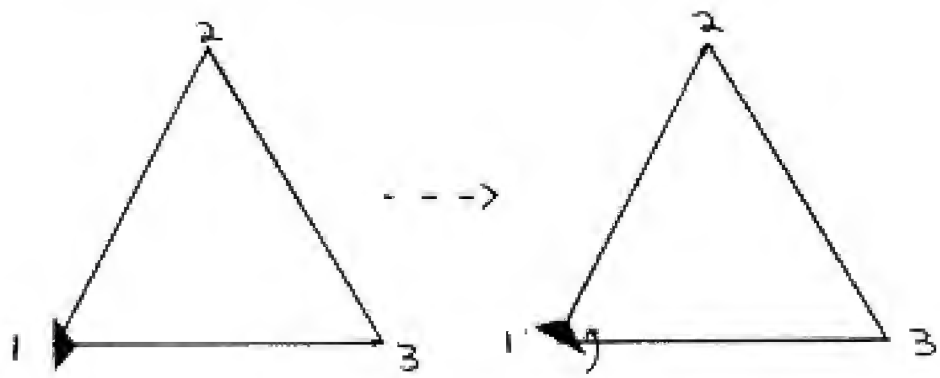

Fig. 1. Measuring angle 1.

He marches to vertex 2 and rotates again, in the same direction as before, to measure angle 2.

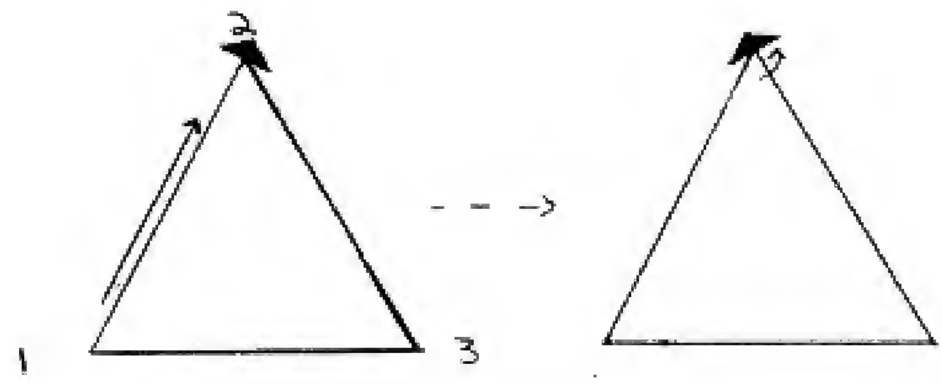

Fig. 2. Measuring angle 2.

(To do this last he must look back over his tail while rotating.) Then he moves to vertex 3 and does the same thing he did at vertex 1. Finally he just returns to where he started, as turtles usually do.

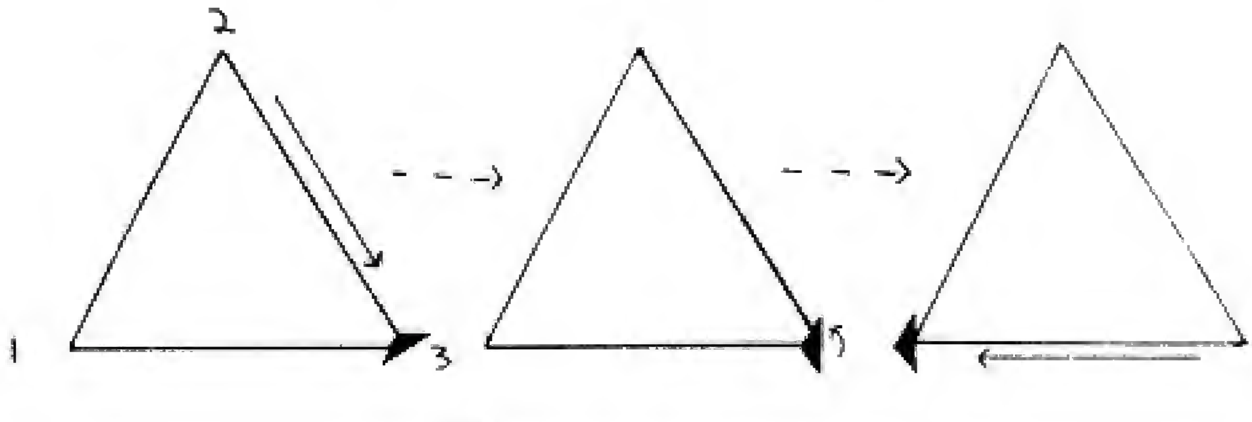

Fig. 3. Measuring angle 3 and going home.

But now look, the turtle has turned through each vertex in the same direction, so his full rotation is just the sum of the vertex angles. And he's now pointing in the direction from 3 to 1 whereas he started out pointing from 1 to 3. In all he has turned exactly 180° ! We don't even have to ask Turtle to remember the separate angles that he measured. Notice that the

final heading of the turtle (180° from initial) does not depend at all on the details of the triangle. Stretch side 1-3, pull vertex 2 off into the distance, the turtle still must end up 180° from his beginning, because having started at 1 pointing at 3, he ends at the same point, pointing along the same line, but in the direction exactly opposite his initial direction.

That sounds pretty solid. After all it proves something we all know (don't we?!) is true. But let me confuse the issue by asking a very hard question. What happens if my turtle is drawing big triangles on the earth rather than little ones on a table top. Is the 180° theorem still true?

Well, you might say, that's just a huge case of a little triangle; the same thing only a million times bigger sides. And since Turtle's triangle measuring process doesn't depend at all on size, he should still find 180°. But careful, watch this!

Suppose Turtle starts on the equator and goes straight north until he gets to (where else?) the North Pole. Now he turns 90° and goes straight south until he gets to the equator. Now he turns 90° and runs along the equator to get back where he started. He's made a triangle. But look carefully, the triangle has 3 90° angles in it! That's 270°. Try that out on a globe.

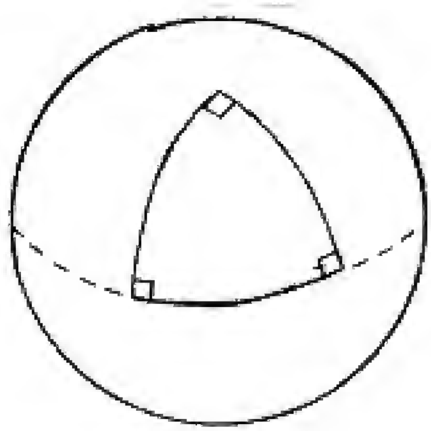

Fig. 4. Triangle with 3 90° vertices.

So what, you say. That's no triangle. Everyone knows a triangle is made up of straight lines, and anybody can see those lines in the "triangle" on the earth are curved. Well, I say, almost anyone can see that, but turtle can't. He's very nearsighted and can't see the curved horizon; he only sees which direction he's going. As far as he's concerned all those "lines" are straight. He's just walking along (like a car with wheels welded straight) not doing any turning at all between equator and pole.

So what, you say again. (You are obstinate.) Turtle, then, just doesn't know enough to know he's not drawing a "real" triangle. But hold on. Let's get some things straight before deciding exactly what's what.

II. TURTLE LINES

First of all we have a new kind of "straight line" to look at, a "turtle line." Turtle can walk on lots of things: the earth, a table top... and whenever he walks without turning, he walks along a "turtle line." (That's a definition, don't argue.)

How about a ping pong ball. Obviously an average size turtle can't walk on it, but I think there are clearly some turtle lines on it. Imagine that the ping pong ball is a little globe, i.e. a map of the earth. You can draw on the equator, and I certainly would want to call that a turtle line. Why? Because a line has the property that if you make a bigger or smaller model of it, it's still a line. So if you shrink the earth with its equator turtle line down to the size of a ping pong ball, I'd still want to call the equator a turtle line. If you want a different reason, imagine a miniature turtle on the ping pong ball; his non-turning paths I'll still call "turtle lines." And I'm sure you'll agree that the "equator" of a ping pong ball must be a tiny turtle line. To find a turtle line,then, all I need is to make sure my turtle is appropriately sized for what he's walking on. In any case he must not be so big that he can't walk around comfortably.

I have a good question. If the equator is a turtle line, is any line of latitude also one? Well, you might initially be inclined to say yes. After all it looks pretty much like the equator. But is it really?

To decide that you have to decide if Turtle can walk along it without turning. Imagine Turtle's little legs churning away. How does Turtle know he's walking in a straight line without looking (he's nearsighted, remember)? To answer that, start out thinking of a simple situation--a table top. I'd also suggest you think how Turtle turns.

Here's my answer: If his left legs take the same number of steps and the same length steps as his right legs he'll go in a "straight line". If he starts taking fewer steps with his right legs (or even takes negative steps) he'll turn. Do you agree? I hope so.

So now, is any latitude a turtle line? Well, Turtle straddles a (Northern Hemisphere) latitude and starts walking. His "south" legs travel on a latitude a bit below and "north" legs a bit above. Marching all the way around the earth has he taken the same number of steps with north as with south legs? Of course not! The more north the latitude the smaller the round trip path. Take a look on a globe again. The equator is the longest latitude, and as you get closer and closer to the North Pole the latitudes get smaller and smaller and eventually reach zero length at the North Pole!

So the turtle must take a different number of steps with his left and right legs, and therefore a latitude is not a turtle line. Do think about that if you're not convinced.

Here's a more clever way to prove a latitude is not a turtle line. It's a symmetry argument, and it works like this. There's one thing we know for sure about turtle lines. There is nothing in the line that distinguishes right from left as the turtle walks along. If he turned, of course he would be turning either right or left; but if he doesn't then there should be no difference in what is right and what is left. Add that to the fact that a sphere does not distinguish one place from another, and you must conclude that on a sphere the left part of the world looking from a turtle line must look exactly like the right part of the world. That is true for the equator or any longitude, but not for any old line of

latitude. Those divide the world into a polar cap and a bigger part, so distinguish right from left and therefore can't be turtle lines. (If you're not convinced by this argument, don't worry. Although I think it's true, I'm not sure it's convincing.)

Turtle lines exist on any surface, it doesn't have to be a flat one or a sphere. Any old bent up surface will do, just set a turtle on it with legs a-churning and watch a turtle line!

But now back to a sphere, the earth. We have a nice triangle made up of turtle lines and 3 90° angles. What happened to the turtle proof that triangles have only 180° ? Something obviously went wrong. As far as Turtle is concerned there are only three places he has to turn (i.e. three places that his right legs don't match the speed of his left legs). And each of those is a 90° angle. But he still winds up pointing 180° from his start. The problem is, if you wind up pointing in the opposite direction on a sphere, it isn't necessarily so that you made 180° worth of turns. Try this! (Get your globe out again.) Start Turtle at the North Pole, notice which way he's pointing. Now walk him straight ahead until he gets to the South Pole. Now don't let him turn at all but walk him sideways (a good straight turtle walk but sideways) clear back up to the North Pole. Presto, he's facing exactly the other way. Turned 180° without "turning" at all. That's the problem with spheres; you can get turned even if you're not turning. If you look from the side you can easily see the sphere turning the turtle without his knowing it.

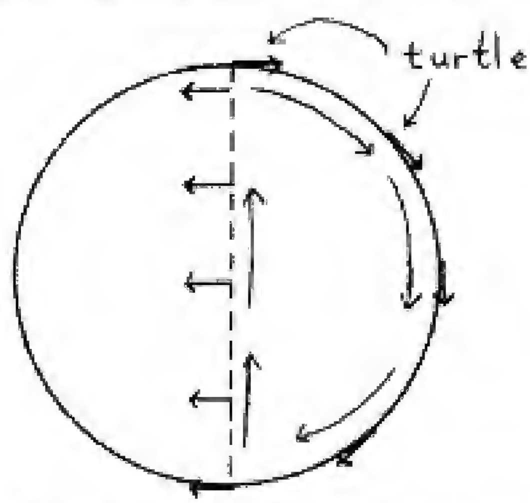

Fig. 5. "Getting turned" without "turtle turning."

Meanwhile back at the 90° 90° 90° earth triangle, it's easy to understand what's wrong with Turtle's original triangle proof of 180°. For sure, from beginning to end of the trip around measuring angles, the turtle has changed his heading by 180° . Evidently the turtle himself turned 270° (3 vertices of 90° each). But now that we realize spheres can turn Turtle without him knowing it, we can hypothesize about that extra 90°. While he was turning 270°, Turtle must have been turned back 90° by the sphere. Turtle turns 270° but

the trip around the sphere turns him back 90° of that.

So you see, there are two kinds of turning: actual "turtle turning" and "trip turning." Turtle only thinks he's turning when he's "turtle turning," but he can be turned by going on a trip even without "turtle turning."

Let's try that idea out (I hope you still have your globe around). We can take Turtle around the triangle without him turning at all by having him walk sideways sometimes. If his trip around the sphere really accounts for that missing 90° (missing between the triangle theorem and the real 270°), Turtle should come back to where he was turned 90° and turned in the direction opposite to the turning he would ordinarily do at the triangle vertices. So start him out on the equator facing East. Walk him sideways up to the North Pole, now walk him straight down (following his nose) back to the equator. When he gets back to the equator he's pointing South and continues pointing South while walking sideways back to his starting point. There he is. Without turning, but merely going on a trip, he's been "turned" 90° . And 90° oppositely from the turns he would make to measure the angles in the triangle. Hooray! The extra 90° in the triangle come from the trip and get added to the 180° of the "triangle theorem."

III. ANGLE EXCESS

Believe it or not, we've made some real mathematical progress in understanding because we have run across a new concept. The concept is what mathematicians call "angle excess" or simply "excess." Excess is the "trip turning" that the turtle gets turned in traveling around a closed path without doing any "turtle turning" of his own on the way. For triangles the excess is exactly the angle you have to add to 180° to find the actual sum of angles in that triangle. Right away there are some nice things to notice about angle excess.

THINGS TO NOTICE:

1) You can ask what it is for any polygon, not just a triangle. (This is provided of course the turtle knows how to walk a straight line in any direction, not just forward or sideways. It is not hard to train turtles to do this.)

2) You can ask about angle excess on any surface, not just a sphere. Simply have the turtle walk around on the surface. So excess is a rather general concept. It's an angle assigned to any closed path on any surface in a particular way.

Perhaps the best thing about it is not its generality but all the nice questions you can ask about it.

THINGS TO WONDER ABOUT:

1) Can you ever compute angle excess without just measuring it? A partial answer you probably guessed already--yes indeed, an excess angle on a plane is always zero! If you still believe the turtle triangle theorem on a plane then you know the excess angle for all triangles is zero in a plane. Even if you can't prove it I bet you'd believe that all polygons in a plane have zero excess. This leads me to ask if the plane is the only surface with zero excess for all polygons? (Think about that.)

2) Is this angle always greater than zero for any surface, i.e. is it always a turn

opposite the triangle measuring direction?

3) In general what does knowing excess tell you about a surface? Everything?

4) What does angle excess really mean?

But let's not get too far ahead of ourselves. I think that we'd better make sure the excess angle concept is nailed down. That means asking some simple questions about it. Is it well-defined, i.e. have we really specified exactly one number to be associated with any polygon? In particular:

Question 1) Does excess angle depend on the initial direction of the turtle faces?

Question 2) Does the it depend on where you start in the polygon?

If you think you have the concept nailed down, skip this section for now, but I'd suggest you come back to it.

Answer 1) No. Imagine the turtle walking around a path without turning and comparing his final heading with his initial to find the excess. Now suppose we take another turtle and start him out pointing 90° away from the first turtle and walk him sideways along the first leg of his trek around the path. Compare the new turtle's heading with the old one's at each point as he goes along. I don't know about you, but I'd say the turtle was broken or didn't know how to walk sideways properly if he gradually changed relative heading to the first turtle, 90° to 89° to 85° to... Somehow the turtle is turning his body to face in a new direction while walking. (Try this idea out by thinking about turtles marching on a table top.) So if the second turtle walks sideways properly, he will maintain his relative positioning to the first turtle and will wind up after going all the way around still 90° different from turtle 1. The difference between initial and final heading must then be the same for the two turtles.

Actually I can be clever and make the argument above into a real proof. I will define a turtle marching sideways (or at any other angle) on a turtle line to be marching "properly" only if his relative heading to a straight marching turtle doesn't change. Then Answer 1) is trivial provided we do all measuring with properly marching turtles. Isn't that clever? This is an example of a great mathematical trick used all the time. If you've got a good theorem and can't prove it, then define things so that it's true.

Answer 2) No. Look at the following record of how a turtle faced as he measured the excess in a triangle starting at A. The excess is marked θ.

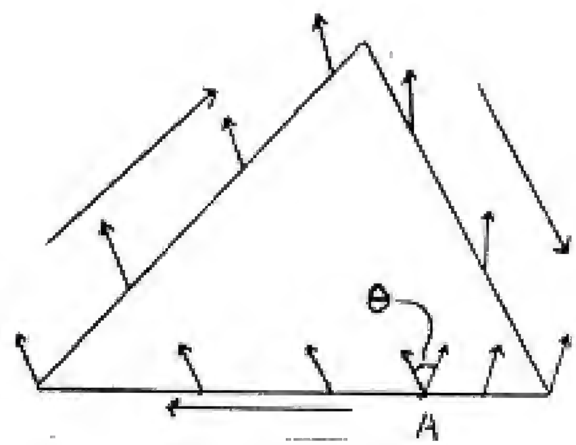

Fig. 6. Measuring Excess.

Suppose some other turtle measures the angle excess starting at B. Because of Answer 1) we might as well take the new turtle to start facing the same way as the old turtle, and so he will face the same way all the way round to A. There the record of the old turtle changes by θ. Answer 1) says that the angle between old and new turtle will be maintained as new moves from A to B.

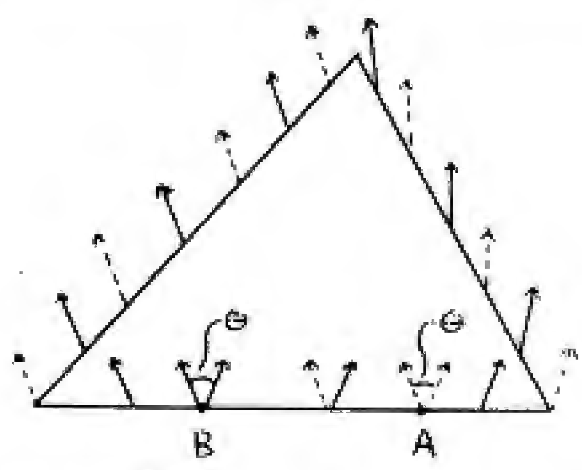

Fig. 7. New Turtle (solid) compared to Old Turtle (dotted).

But then θ will be the angle between the beginning and the end of new turtle's trip, the same as old turtle's excess angle.

The result of Answer 1) and Answer 2) is that it doesn't matter where you start or at what heading you start, the excess will be the same. (Note that this had better be true, otherwise the understanding of excess as being the difference between triangle vertex sum and 180° would likely be in error. Capisce? (This is more or less obvious depending on how you think about it. In any case you should think about it until it's more or less obvious.))

IV. EXCESS ADDS!

Back to more investigation and less formalization. Let's concentrate on the sphere for awhile. I started on the sphere with a triangle of fairly large excess, 90°. Can you imagine a triangle with a bigger excess? (Look for one if you have time.)

How about this one. A triangle with 3 180° vertices! The equator! Just start anywhere, call it point 1. Travel 1/3 of the way around then stick in a 180° vertex there and continue around on the equator another 1/3 circumference to point 3. Stick in another 180° vertex and run the rest of the way round home to 1. An excess of 360° ! That's one of my favorite triangles.

How about a triangle with a smaller excess. I guess that's no problem. Just take a very small triangle, like one in your front yard, compared to the earth-sized sphere (the earth!). I'd bet such a triangle has angles totalling very nearly 180° so an excess of nearly 0.

It looks like small triangles have small excesses and large ones have large excesses. If you haven't already noticed, the 3 X 180° triangle has 4 times the excess of the 3 X 90° triangle, and can be made up by pasting together exactly 4 of the 3 X 90° triangles. (Look closely.) From that it looks like excesses add. Take a look at a triangle made by combining 2, 3 X 90° triangles. It has a excess of 180°! Look at any triangle made up of 2 90° angles at the equator and n degrees at the pole. It has excess of n° and can be made up of n triangles of one degree excess.

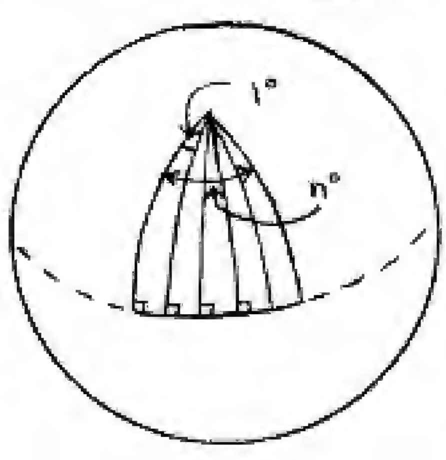

Fig. 8. Excess of the large triangle is sum of the excesses of the small ones.

This is beginning to look like a theorem.

> Theorem: If a triangle is subdivided into subtriangles then the excess in the triangle is the sum of the excesses in the subdivisions.

I'll give you a proof so you don't have to fiddle around a lot. Notice that the theorem doesn't mention anything about spheres in particular, neither will the proof; it's true on any surface.

I'll just do the case of subdivisions into two. It's tricky (really!) extending this to any subdivision, but I think just this much should give you an idea of what's going on.

Proof: Measure excess in ABC. I've shown a record of turtle pointing.

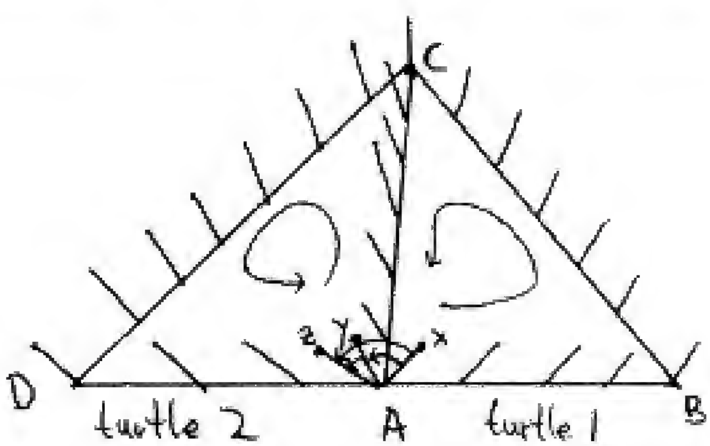

Fig. 9. Excess BCD (x to z) = Excess ABC (x to y) + Excess ACD (y to z).

He starts out with heading x and ends up with y. Now do ACD. To make things simple you might as well start from A with the same heading, y, which ended measuring ABC. Then the second measuring turtle should agree with the first all the way to C. (Do you agree?) Turtle 2 then continues to D and ends up at A with heading z. Finally measure the excess in BCD by starting at A in position x, running all turtle 1's path to C. Then follow turtle 2 (who starts from C with the same heading turtle 1 left off with) around clear back to A. He winds up with heading z. Look, the excess of the big triangle is the angle x to z which is just the sum of the excesses (x to y and y to z) of the two smaller triangles.

That's really a pretty nice theorem. It's the beginning of a really great one.

> Theorem: The excess of any polygon is the sum of the excesses in any

polygonal subdivision.

Can you see how to prove that? It's not that important, but it's nice to notice that all you have to do is start with one polygon and add on connected ones, one at a time. To prove each step in that process you can use the proof given for adding two triangles together because that proof did not depend on the pieces being triangles! All it needs is a picture like below. Polygon vertices are irrelevent.

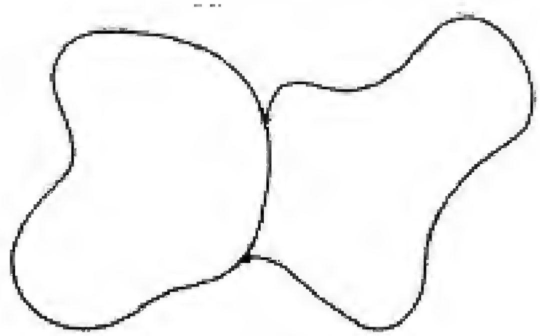

Fig. 10. The Topology of Fig 9.

In mathematical lingo, the proof only depends on the topology i.e. how the thing is hooked together, who's connected to whom. Not on where vertices are or how long any side is or how much area anybody has.

V. CURVATURE DENSITY -- EXCESS PER UNIT AREA

Whether or not you spent time proving that theorem, it's certainly suggested by our observations about excess being additive in some special cases. And furthermore the theorem is really suggestive of other things. Compare it to the obvious:

> Theorem: The area of any polygon is the sum of the areas of any polygonal subdivision.

Well, excess acts like area in that respect. Could it be that excess is proportional to area, that is $E=kA$ where E=excess A=area and k=some constant? That would account for the additivity of excess. But it's obvious that k couldn't be a universal constant like π. After all, k must be zero for a plane, but it can't be zero for a sphere. Not only that but it can't even be the same constant for all spheres. Consider also a $3 \times 90^\circ$ triangle on the earth and

one on a ping pong ball. They have the same excess but certainly don't have the same area. So how about the hypothesis that k depends on the surface, and every surface may have a different k? Let's try that out.

Obviously any plane has $k=0$.

> Theorem: On a sphere of radius r, $E=kA$ and $k=1/r^2$ (if E is measured in radians).

Proof: This is not a rigorous proof just as the other proofs I've given are not rigorous. But the main idea is there. First I want to show $E=kA$. Then it will be simple to find out what k is.

The idea is to measure excess as you measure area. Subdivide a general area into a bunch of tiny uniform areas and add them all up. For example how might you measure the area of an arbitrary polygon. Divide it into little tiny squares and count squares. (Of course there can be a little area left over but not much. I can always use smaller squares to get a better estimate if I want. And you calculus buffs out there know how to talk about the limit of small squares.)

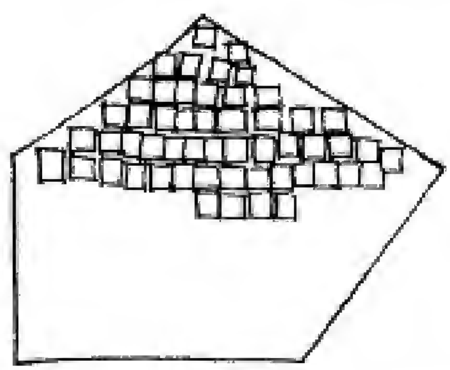

Fig. 11. Dividing a polygon into small squares -- measuring area.

So do the same for excess. Divide your polygon up into little tiny uniform "squares" (again there may be some area left over, even cracks between "squares", but again if your "squares" are small enough--not much will be left over). Now these "squares" are all absolutely identical so not only must they have the same area, they must have the same excess. If the area of a small "square" is a, its excess is e and there are N squares, then $A=Na$ and $E=Ne$ and thus $E=(e/a)A$. If you take another polygon, you can measure its A and E using the same small squares and if this new polygon has M squares $A=Ma$ and $E=Me$, still we have $E=(e/a)A$. The same e and a so same $k=e/a$.

So $k=e/a = E/A$. We still need to find what number k is for a sphere. Well, take an example. For the 3 x 180° triangle $A=1/2$ sphere area $=2\pi r^2$. $E=2\pi$ (360°). $k=E/A = 2\pi/2\pi r^2 = 1/r^2$.

There are a few things to be said here. The above proof that $E=kA$ applies to any surface which is the same everywhere. It applies to a plane ($k=0$) to a sphere ($k=1/r^2$) and to anything else which is the same everywhere. By "the same everywhere" I mean that you can use the same small square as your little measuring stick everywhere. Certainly you can move any little square on a sphere to any other place on the sphere. The same in a plane. But suppose that you have a sphere and drop it so that one side gets flattened. Then you just can't move a little square from the round part of your smashed sphere to the flat part to measure some A and E. It won't fit. After all flat is flat and round is rounded. You must use a different "tiny square" reference for measuring E and A on the sphere part and on the flat part of your smashed sphere. So $e/a=k$ is different on the two parts.

I said that $k=0$ for a plane which is not curved at all. A very large sphere like the earth doesn't look very curved and indeed a little chunk of sphere like your back yard (or part of a very calm lake if your back yard is too rocky for your taste) could easily be mistaken for a flat plane. k is very small in this situation since $k=1/r^2$ and r is big. Now a ping pong ball is very curved compared to a large sphere hence $k=1/r^2$ is very large.

Let me make a very interesting analogy. $k=e/a$ is a "density". It's like how much paint you have per little chunk of surface area. In fact I'll call k the "curvature density." and I want to think of it as analagous to "paint density." Spheres and planes have uniform coats of "paint"—that is the curvature density is the same at all places. But just like a room which has perhaps 1 coat of paint on the wall (1/2 cup per sq. foot) and a double coat on the woodwork (1 cup per square foot) and no paint at all on the windows, the curvature density may vary from place to place on a surface. The flattened sphere has no curvature density on the flat side and $1/r^2$ on the side that hasn't been smashed. A football is not too curved in the middle, k is not too big there. But at the pointed ends a football is curved as much as a sphere of small radius. It has a thick coat of paint--I mean a large curvature density there. Between the middle and the pointed ends the "paint" probably gets gradually thicker and thicker.

If you want to know how much total paint, P, on some surface with constant paint per unit area, p, the answer is simple, $P=pA$, where A is the total area. In the same way the total curvature, K, of a surface of constant curvature density is just $K=kA$. Now if the paint density varies from place to place and you want to know how much total paint there is, how do you do it? My answer is divide your surface into little tiny pieces. You find out how much paint is on each little piece by multiplying paint density times area, and then add to find the total paint. (If you have only have two thicknesses of paint then you only have to divide your surface into two pieces, take density times area for each one. But if you have lots of different thicknesses then you're probably best off dividing into lots of small areas.) In any case I imagine you believe that if you know the paint or curvature density everywhere you can figure out how much total paint or total curvature is on the surface.

ESCAPADE: I want to look at curvature on another surface, a cylinder.

Q: What is k on a cylinder?

A: k=0 like a plane! Obviously a cylinder is "the same everywhere" so k is just a single number. I'm sure you'd agree that lines running the length of the cylinder and circles perpendicular to those lines are turtle lines.

So look at the square below.

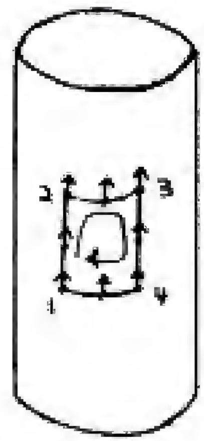

Fig. 12. "Square" on a cylinder.

Turtle goes forward 1 to 2 sideways 2 to 3 backwards 3 to 4 and sideways 4 to 1, and he winds up not turned at all. E=0 so if E=kA then k=0.

Now why is that? The reason it's true is profound. It's that a cylinder is just a plane rolled up. And rolling something up doesn't change any path lengths on the surface. (Demonstration: draw a path on a piece of paper, now change that path length. Well, you can't, not without ripping the paper. So rolling doesn't change any lengths. For those who are not easily convinced, I suggest the following. Glue a rubber band to a piece of paper. Now try to stretch the rubber band without ripping the paper. Rolling just won't do it (except for a tiny bit which happens because the rubber band is not on the surface but a little above.)

If path lengths don't change on rolling, then a straight (turtle) line drawn on a paper which is then rolled up, remains a turtle line. How can I be so convinced that straight lines don't become non-turtle lines? Because having thought about how Turtle runs along turtle lines I know that all it depends on is that a bunch of distances are equal. Turtle knows he is walking in a straight line when his left legs and his right legs are moving the same distance in each step.

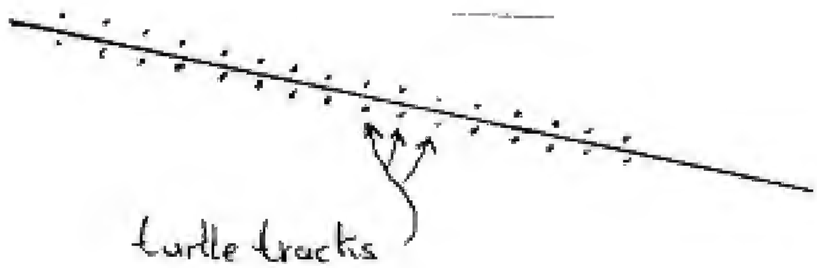

Fig. 13. A turtle line with tracks.

Look at the above straight line and turtle tracks around it. Now imagine the paper rolled up. Turtle could walk in the exact same tracks because no distances have changed. A turtle line remains a turtle line when rolled up.

There's more. Angles don't change when you roll something up or unroll it either. So any polygon of turtle lines on a cylinder is the just the same as a polygon in a plane. It has zero excess.

If angles and distances don't change when you roll a paper into a cylinder, what does? Topology! Here it is again. The only thing that changes is who is connected to whom.

There's a rather important lesson here. When we began with planes and spheres, it was pretty clear that k=0 meant what people usually mean by saying "flat" and k not equal to zero meant curved. But now here's a surface, the cylinder, which most people would say is curved. You have to decide now whether you want to go on saying a cylinder is curved as you always did, or change your definition to "flat" means k=0, and then say that a cylinder is flat. That last may sound very strange, but mathematicians (and I) think it's a good thing to do. Why? Because we are interested in geometry like how many degrees in a triangle and squares having right angles etc. We are not interested so much in "how things look." A plane has so much more in common with a cylinder from a geometric point of view than a cylinder has with a sphere, that it makes much more sense to say both a plane and a cylinder are flat (rather than saying spheres and cylinders are curved). In fact if a turtle were never allowed to go clear around the cylinder and discover its different topology, he'd never be able to tell the difference between that cylinder and a plane at all! So if you're talking to your friends who haven't read this paper you'd be better off saying a cylinder is curved, but if you're talking to a mathematician, he'd be happier to hear you say a cylinder is flat. If you're doing geometry it's hard to go wrong with flat cylinders.

VI. REVIEW

Let's stop and catch our breath. Take a look at what we've done. We started with your usual garden variety straight line and things you can build out of them, like triangles. You can ask yourself what really is a straight line and there are lots of ways to answer that. One useful way is with a turtle walk. If a turtle walks an equal number of steps with right and left legs and equal distance in a step, then that's a straight line. But that way of characterizing a line works just as well on a sphere or a football or a cylinder as it does on a plane. The question arises, what happens to things like triangles with these "turtle lines" for sides. The main thing is that certain angles or sums of angles change. We found that you can think of this change as being described by a new angle, the angle rotation which the surface performs on a turtle (as opposed to that which a turtle does himself) in travelling around a polygon. That rotation is called the excess. Then there was our prime theorem about excess. It is additive: the excess of a polygon which is subdivided is just the sum of the excesss of the subdivisions. That makes excess look very much like area, and in fact for surfaces which are everywhere the same, excess is just proportional to area. In other surfaces the amount of excess per unit area varies from place to place being greatest where the thing is curved most and less where it is curved very little. That all leads us to define curvature density, k, as the excess per unit area in a particular place on a surface. (It's $k=1/r^2$ on the rounded part of the dropped sphere and $k=0$ on the flat part. Don't ask what it is on the edge between rounded and flat just yet.)

You may ask me why I called $k=e/a$ the curvature density rather than excess density, after all you compute $k=e/a$ by measuring how much excess occurs per unit area. Or why did I call $K=kA$ the total curvature rather than the total excess? Well, one reason is that while k is measured using angle excess, it is really important for its "meaning" as how curved something is. It's always good to remind yourself of what something means in its name. There is another reason which is really quite subtle though. I know how to compute the total curvature on any kind of surface by adding up curvature density times area of little pieces. For example how about the following surface.

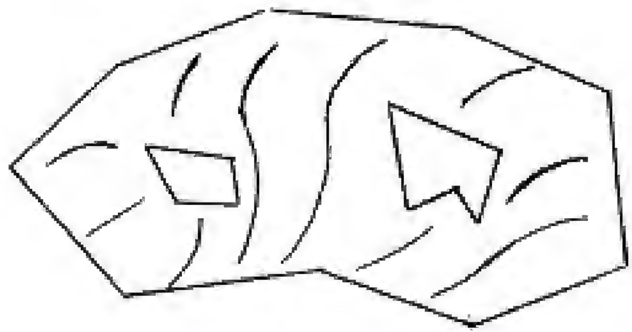

Fig. 14. An arbitrary surface with holes.

Suppose, to be simple minded, that e/a=constant. No problem you say (I hope) if you can measure area. But if I try to think of e/a as density of excess, then I may be tempted to think of the total curvature as a total excess. But where is that excess? An excess always belongs to a closed path. And I'll give you a hint, the total curvature of the above shaped surface will not be the excess angle of any of its three edges!

Here's a different example of the same thing. What's the total curvature of a sphere? Well $k=1/r^2$ and $A=4\pi r^2$ so $K=4\pi$. That sounds like an excess angle, but where is the path for that excess??? That's why I use the names "curvature density" for e/a and "total curvature" for K=kA rather than anything involving excess which may make you start looking for a path.

Let me go back to the holes business and look a bit more carefully. Suppose you draw a polygon on an arbitrary surface, and you want to measure the total curvature inside it. So you divide it up into lots of tiny polygons and add up k times a for all the small pieces. You're just adding up the e from each small polygon. But doesn't the additivity theorem say that that sum is just E of the big polygon! Look at an example. Suppose that the surface we are talking about is a sphere with the polar cap cut out (say from latitude 80 on up). Take the equator as a polygon, a polygon with one side! You know it has an angle excess of 2π (360°). But the total curvature contained in the upper part of the sphere is ($k=e/a=1/r^2$) times (an area less than $2\pi r^2$). Thus the total curvature inside cannot be the excess of the equator. Something is rotten in the state of the additivity theorem!

Let me tell you what's gone wrong so you don't have to figure it out. The additivity theorem needs something that I didn't mention explicitly. It needs the polygon to have an inside which topologically looks just like the inside of a polygon in a plane. That means no holes, no tears or other such gobbeldy-gook. (Those holes and tears won't allow the proof by adding pieces one at a time to work. That's because we can only add together

pieces which have exactly the topology of Figure 10. With a hole you always come to a situation where adding on a piece does not have the Figure 10 topology. Try working out an explicit example!)

By the way, holes aren't the only things that can mess up the additivity theorem. The "square with handle" shown in Figure 29 has no holes, only one edge (the square), yet the excess of the square will not be equal to the total curvature inside it. We'll learn more about that kind of thing later.

Just to remind you that we're still doing mathematics I'll restate the above discussion in a theorem and give a proof. The reason I'm doing this extra work is because this is really a key theorem and an example of a class of theorems which are very important in math and physics. It's a theorem which relates something that depends on the interior of a region to something which can be computed on the boundary of that region. You'll see how such theorems can be important soon.

> Theorem: Given a polygon on an arbitrary surface which has an interior topologically like the interior of a polygon in a plane, the angle excess of the polygon is equal to the total curvature of its interior.

Proof: Let's compute the total curvature in the interior. To do that divide the polygon into lots of very tiny ones. Then multiply the curvature density in each small polygon by its area and add up the results. But the curvature density, k, of each tiny polygon is just e/a and so k times the area is e. We are really just adding up the excesses of the small polygons. The additivity theorem now says that the thing we're computing, total curvature, is just the excess of the large polygon. QED.

This theorem tells you exactly when $E = K$.

Problem: What's the total curvature of a sphere? Previous Answer: A sphere is the same everywhere so $K=kA$. $k=1/r^2$, $A=4\pi r^2$ so $K=4\pi$. Another Answer: to find total curvature we can add up the curvature from any convenient pieces. A sphere is its Northern Hemisphere plus its Southern Hemisphere. Each of these is a polygon (the equator) bounding a nice interior. So the curvature of each hemisphere is just the angle excess of its boundary, the equator. We know the equator has excess 2π. That makes the curvature for each hemisphere 2π, and a total of 4π for the sphere.

VII. DENTS AND BENDS

Suppose I put a little dent in a sphere, what happens to the total curvature? You might guess lots of things. Perhaps it depends on the dent. If it's a flattening maybe that reduces it; if it's a pointy kind of outward dent maybe that increases it. But the answer is <u>nothing happens to the total curvature!!!!</u> Watch carefully how I prove this.

Suppose you make a dent in the sphere. Let me draw a polygon around the dent but far enough from it so that the vicinity of the polygon is unaffected by the dent.

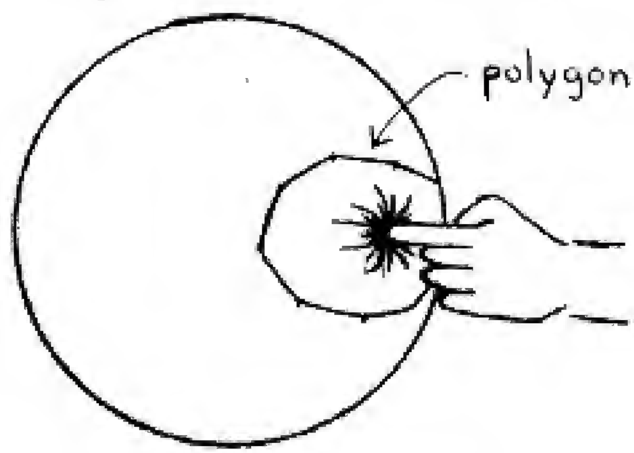

Fig. 15. Making a dent in a sphere.

The total curvature of the sphere is just that inside the polygon plus that outside. The total curvature outside the polygon cannot be affected by the dent so all we have to worry about is the total curvature inside. But that's just equal to the excess of the polygon. And the turtle's walk around the polygon is entirely unaffected by the dent since I drew the polygon deliberately far enough away from the dent so that there's no bending there. (The turtle before and after denting treks the same territory.) The excess must be the same before and after denting, and the total curvature does not change at all!

I can dent, bend, smash, buckle, push, pull a sphere one small piece at a time into any shape I choose and the total curvature remains the same. The total curvature is a topological invariant, that is, it doesn't change no matter what you do to the sphere as long as you don't rip it or in some other way change its topology. I think that is really marvelous. So a football has total curvature 4π. A Mickey Mouse balloon, ears and all has total curvature 4π. That sphere I dropped earlier and smashed one side flat still has total curvature 4π. (Whoa. Suppose I smash the Southern hemisphere flat, you say. That part has total curvature zero. The Northern hemisphere has only 2π. Where'd the other curvature go? It's there! Find it!)

Suppose you take a sphere and pull it to make a cylinder capped on each side with a hemisphere.

Fig. 16. Cylinder between two hemispheres.

That thing has total curvature 4π. But each cap has 2π, so that doesn't leave you anything for the cylinder. Cylinders have total curvature zero! And since a cylinder is the same everywhere, the curvature density must be zero everywhere. (Now we've shown that two ways, it must be true.)

All this bending and stretching and moving curvature around makes me ask, what kind of bending and stretching you can do to something, even if it's not a whole sphere, and keep the same total curvature. Pretty clearly I'll keep the same curvature as long as I can put an unbent "turtle road" around the dent and isolate the rest of the surface from the dent. (Well that may not be so obvious as you think, but it's a good working hypothesis.) In any case it is definitely true that if the surface I'm talking about is a turtle polygon with a nice interior (no rips or funny business) then the additivity theorem tells me I can compute the total curvature by running a turtle along the edge. So if I leave a little ribbon undented along the edge of the polygon as a "turtle road," I can be sure I haven't changed the total curvature of a nice polygon.

Can you extend this discussion about maintaining total curvature to surfaces which are not nice polygons but might have holes or rips in them? How about a polygon with an interior which has a "handle?" Can you give an example where you keep the edge polygon of a surface unbent but don't maintain an entire little ribbon "turtle road" and consequently change the total curvature?

Here's another way to keep the same total curvature. Just make a bigger or smaller model of your surface. If your surface is a polygon with a nice interior then total curvature is just an angle (the angle excess). Angles don't change when you change just the scale of something. Can you prove in general that the total curvature of any surface doesn't change on making a bigger model of it? Hint: Think about how you measure total curvature.

VIII. SURGERY

How much total curvature in the surface of a donut—mathematicians call it a torus. The first thing to notice is that any donut has the same total curvature as any other. Why? The same reason spheres, footballs and Mickey Mouse balloons have the same curvature. I can isolate a small dent with a turtle path and show the total curvature doesn't change. Then I can make any number and kind of small dents to bend a donut into the shape of another donut and the curvature stays the same at each step.

I might try to do for a donut the same as I did for a sphere. Unfortunately donuts are not the same everywhere. Inside the hole a little chunk of surface looks like a saddle but outside, the surface looks a lot more like just a cap on a sphere. So I can't use $K=kA$.

What I'll do is start with something I know, a sphere, and do some surgery on it to make a torus. If I can figure out what happens to the curvature during surgery, I win.

So take a sphere. Squash the North and South Poles inward together. Still has curvature 4π.

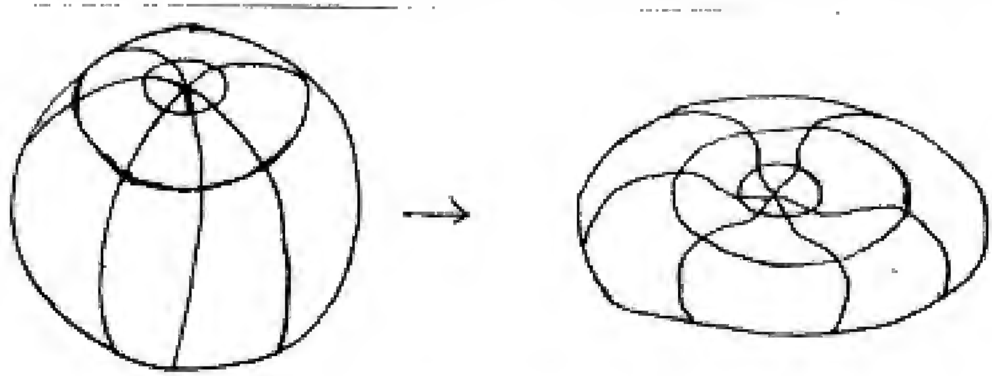

Fig. 17. Squashing a sphere.

Be sure to make the middle where the poles nearly touch flat so that there is no curvature density there. Now cut a small hole out of the North and South Poles.

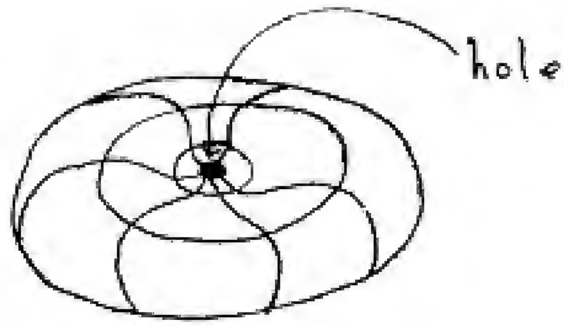

Fig. 18. Removing a small, flat (with zero curvature) disc.

That does not remove any of the curvature. Now just insert a little valve shape like below, to make a donut. (Top edge of valve fits in North Pole hole, and bottom fits in South.)

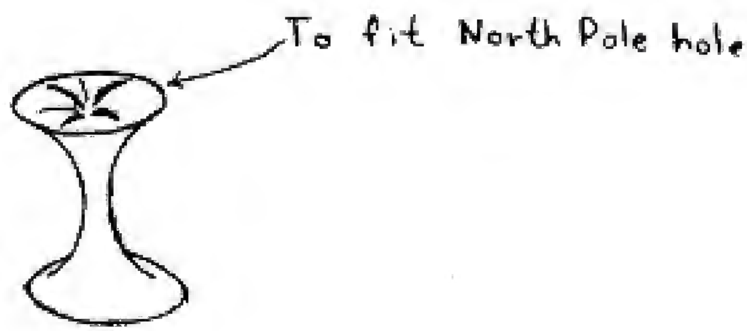

Fig. 19. A valve.

If you like, pull the center hole out until it looks like a real donut.

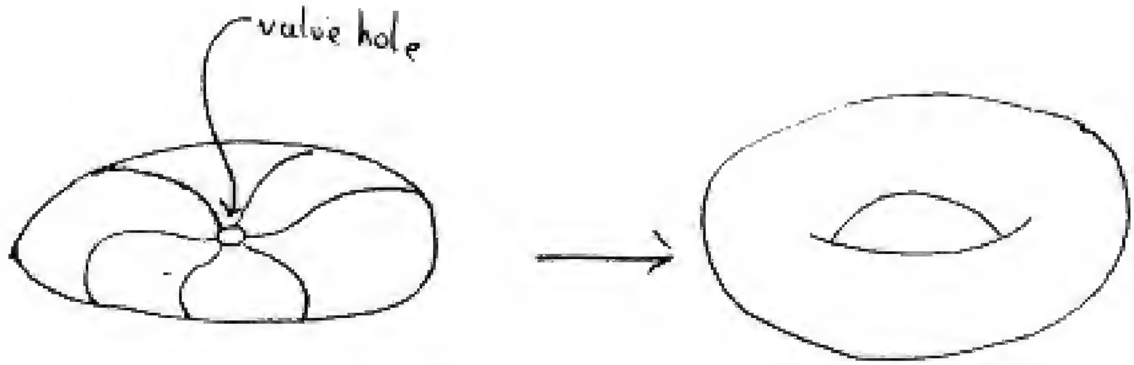

Fig. 20. Stretching the valve hole.

The only place where curvature was added was when we glued on the little valve. All we have to do is figure out how much curvature is in the valve and add that to the 4π from the sphere. If I can construct the valve out of surfaces of known curvature, that would do it. Try this. Start with a sphere (total curvature equals 4π). Put a belt around the equator and squeeze, then straighten things out so that the North Pole bulge and the South Pole bulge are spheres. Total curvature 4π on top bulge, 4π on the bottom bulge and look who's in between.

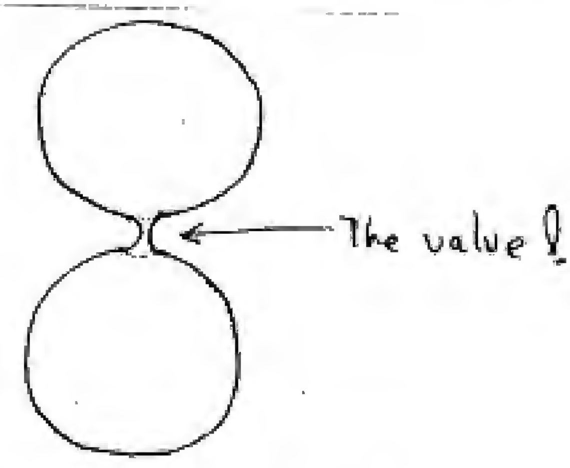

Fig. 21. Barbell = Sphere + Valve + Sphere
Total Curvature: $4\pi = 4\pi$ + Valve + 4π
Valve = -4π.

The valve! He must have curvature minus 4π so that this "barbell" still has total curvature 4π. Thus the donut which we constructed from a sphere and a valve has total curvature 4π minus 4π equals zero. That does not mean a donut has curvature zero everywhere, but just that it has as much negative curvature as positive curvature.

Now as an exercise how much curvature is in a two holed donut?

Fig. 22. A two holed donut.

IX. CONES -- ANOTHER LOOK AT CURVATURE

I'd like to go back and look at cones for a bit. Now a cone is mostly just rolled up flat paper. Cut a little chunk out of the side of the cone and you can easily lay it flat. That means that the cone is just about everywhere a zero curvature density object. But there is an exception, the tip. The tip won't flatten out without ripping. All the curvature in a cone must be concentrated at the tip.

How can you figure out how much curvature is in in the tip? Of course--use a turtle and find out how much excess is contained in some path around the tip. So that I can easily see a turtle walking around on the cone, let me do a little trick. Cut the cone up the side and lay it out.

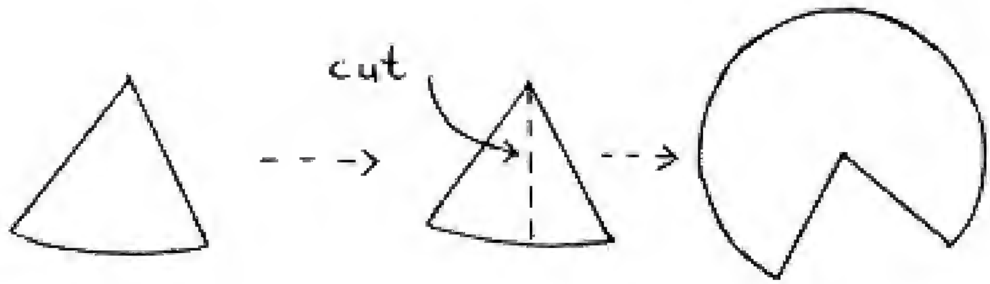

Fig. 23. Laying a cone flat.

This doesn't affect any distances, and turtle paths can easily be seen now; they are regular straight lines (except where they cross the cut). I can also tell easily when the turtle doesn't turn; he just keeps the same heading. So now look at a turtle path around the point. I've drawn in the direction the turtle points along the way. (I started him at A pointing parallel to the cut.) If I glue the cone back together it's easy to see that the excess for the now closed path is exactly the angle θ between the turtle at B and the cut.

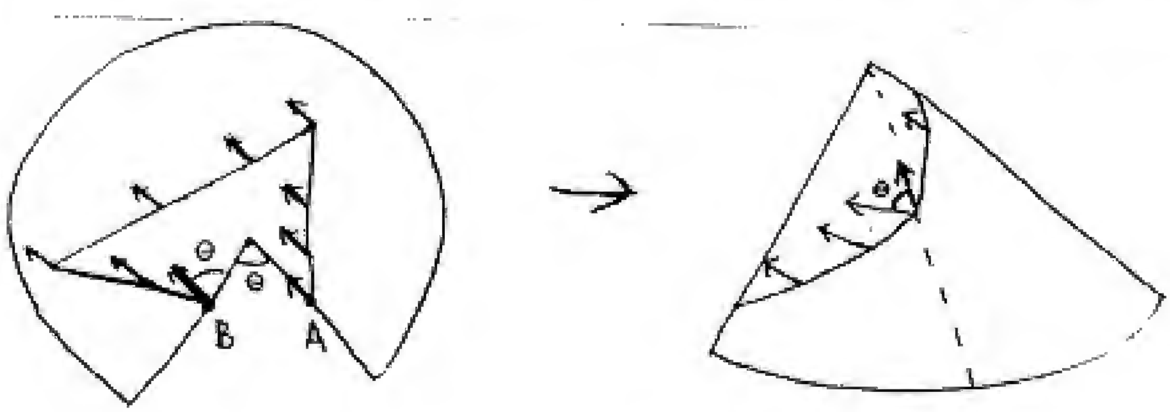

Fig. 24. Turtle path with turtle headings.

But θ is exactly the angle of the pie cut out of the cone when laid flat. (That's just a little

elementary geometry for you.) So the excess of the turtle path is just the "pie angle."

Notice that this result, excess of path around tip equals "pie angle" does not depend at all on how big the path around the tip is. So you can see by pushing the path closer and closer to the tip and always getting the same excess that the curvature must be concentrated in the tip with zero curvature everywhere else.

> What's nice about cones, then, is that you can see the angle excess. It is in fact the angle you need to cut from a flat piece of surface to make it into a cone.

Suppose Turtle is sitting on the apex of a cone. And then he goes a distance r away and draws a "circle" (actually a many-sided polygon) around the apex. Being nearsighted and not too concerned with curvature he thinks r is the radius of his circle. But of course he finds the circumference is not $2\pi r$ but something less. (It's missing exactly the "defect pie" from being $2\pi r$.) He'll find the same thing on a sphere. The circumference of a circle is not $2\pi r$, but a little less. (Look at a globe and what the turtle thinks is r.) In general positively curved things ($k>0$) have this property of circles having "insufficient" circumference. Of course that really is why they are curved; there's insufficient circumference to a circle to allow you to push it flat without ripping the thing.

How about negative curvature? It's like a cone with too much rather than insufficient "circumference." Not a pie with a slice taken out but a pie with an extra slice.

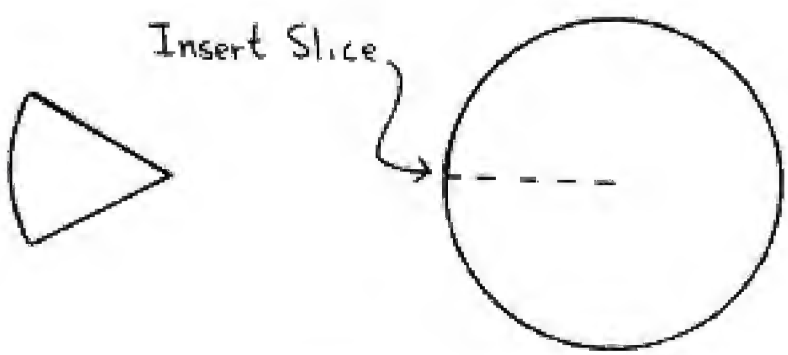

Fig. 25. Negative curvature.

You can't push such a thing flat. Not because you're short of circumference and will rip the cone trying to flatten, but because you have too much circumference for a given radius and can't cram it all into a plane. Saddles have negative curvature.

X. PROBLEMS

1 through n:(Author's prerogative) Answer all the (interesting) questions in the text.

n+1:(Back to basics) Convince yourself that an equator must be a turtle line independent of the fact that I told you it was. Is the path of a boat with rudder aimed straight a turtle line? How about a jet plane flying straight?

n+2:(Cutting and pasting) Is a circle around the apex of a cone a turtle line? Make yourself a cone and draw some turtle lines on it just to see how they look. Draw in some unturning turtle direction flags.

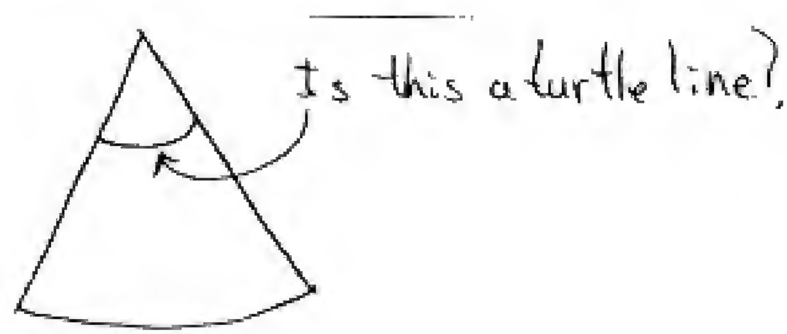

Fig. 26. "Circle" around cone's apex.

n+3:(An easy one) Show that a football has total curvature 4π without using the fact that it is a bent sphere. Hint: Find a subdivision of a football into nice polygons of which you know the excesses.

n+4:(But I already knew that) Using the theorem that making a bigger or smaller model does not change total curvature, conclude that all the curvature in a cone is in its tip. Do this by observing that a smaller model of a cone has the same curvature but is exactly a smaller piece of the original cone.

n+5:(Some idiot poured paint in my garden hose) Sometimes it is not useful to describe the location of paint by paint density (paint per unit area). After all when it's still in the can you just say, "there's a gallon right there." The curvature in a cone is like that. It's all in one place, the tip. On the other hand some crazy person might pour your paint into a garden hose, and then the most reasonable measure would be paint per unit length (of

hose). Can you find a surface where that sort of measure is appropriate for curvature. Hint: The edges of a cube are not an example! Convince yourself that they contain no curvature. (By the way, where is the curvature in a cube if it is not distributed along the edges?)

n+6:(Changing scale) Convince yourself that making a larger model of any surface changes the curvature density by a factor of f^2 (f is the factor of increase of all dimensions from original surface to model), but that the total curvature of model and surface remains the same.

n+7:(Turtle lines revisited) Look at the following record of turtle tracks.

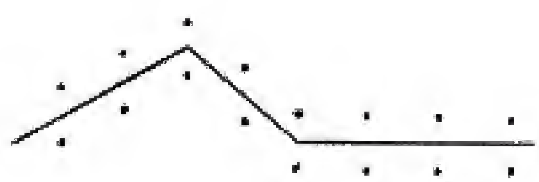

Fig. 27. A turtle line?

The left legs take the same number of steps as the right legs. And all steps are the same length. So why isn't the track a turtle line? (It obviously isn't one.) Can you apply the principle of "a line must be everywhere the same." Can you answer the question without the principle?

Another problem along the same lines is as follows: If a line must be "everywhere the same," then what happens to a turtle line on a smashed sphere as it goes from the round part to the flat part? Can you reformulate the "a line is everywhere the same" principle so that it really applies to turtle lines? Think of a turtle line as a procedure.

n+8:(More turtle lines revisited) Would the little turtle which we used for ping pong balls draw the same turtle lines on a big sphere as a big turtle? What do you have to say about turning a real (motors and gears) turtle loose to draw triangles in your back yard. Would his size matter? Think about a tiny,tiny turtle crawling over each pebble in your back yard. Does that make you nervous about what a turtle line really is? I mean you know pretty

much what a triangle of turtle lines say 20 feet on a side should look like on your back lawn but wouldn't a tiny turtle get all confused by the blades of grass? Who tells you what size turtle to use?

Ask a mathematician to answer this question. Ask a physicist.

n+9:(Relatively speaking) Suppose somebody told you that not only is the earth not flat, our UNIVERSE is not flat either. What might he mean by that? Answer in terms of angles, circumferences and radii of circles, and perhaps surface area and radii of spheres. Notice before starting that we never had to go out of the surface of a sphere to discover it was curved as long as we had turtle lines drawn. The same applies to the universe. (Einstein decided that the universe was not flat. He thought it was curved in a very special way to account for the existence of gravity. In fact he arranged things so that the funny shaped paths objects travel under the influence of gravity are just "turtle paths" in our curved universe.)

n+10:(Holes) Can you find the total curvature of a surface with holes in it in terms of the excesses of the boundaries of the surface? Hint: Fill in the holes. Now how much curvature does the surface have? Cut them out. How much curvature did you remove?

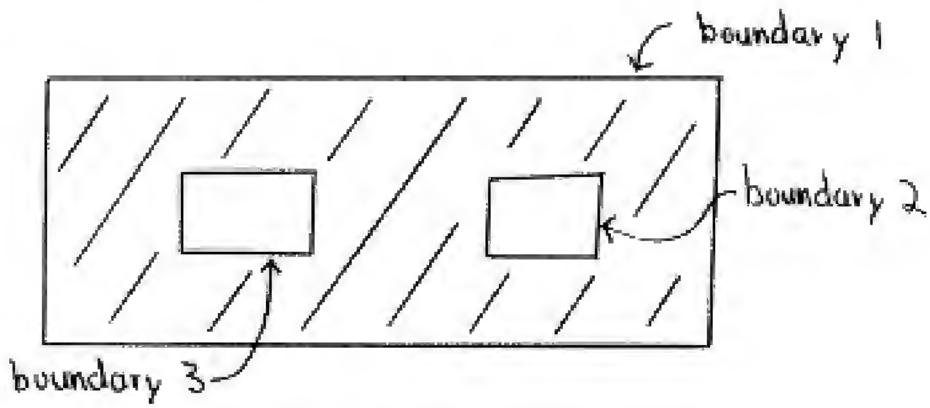

Fig. 28. Surface with two holes.

n+11:(Handles) By doing some bending and surgery show that the addition of a handle to a surface always decreases the total curvature by 4π.

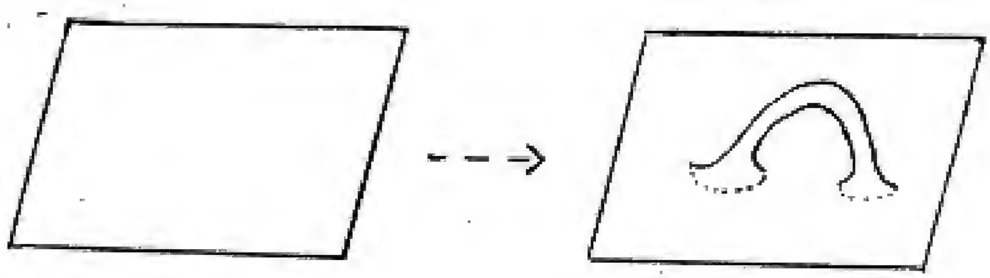

Fig. 29. Inserting a "handle."

(Adding a handle is topologically drilling two holes and gluing in a bent cylinder (valve) to connect them.)

By doing surgery, show that a knotted handle can be untied without changing total curvature.

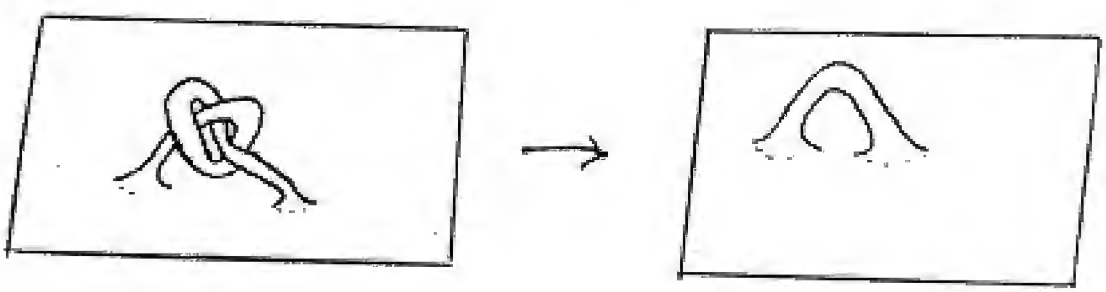

Fig. 30. Untying a knotted "handle."

You may want to use the fact that an unbent cylinder has zero curvature so you can saw out part of it without losing any curvature. Alternatively you can prove and/or use the theorem that bending any surface does not change total curvature as long as all edges

remain unbent.

Finally as an interesting point for you, consider the class of surfaces which can be made from a limited amount of area (thus planes are excluded) and have no edges. A sphere and a torus are examples. Now the remarkable fact is that, topologically speaking, every such surface is just a sphere with handles stuck in. (A torus is topologically a sphere with one handle.) The preceeding discussion of handles should convince you that every such surface has total curvature belonging to the set 4π, 0 -4π, -8π, ...etc. and that the total curvature tells you exactly how many handles the surface has. (Note: This little discussion refers only to the garden variety surfaces found in ordinary three dimensional space.)

n+12:(Plato) A Platonic solid is any object which is flat almost everywhere and otherwise is as "regular" as can be. That means its surface is made up of a number of faces which are all identical regular polygons pasted together. (A regular polygon has all sides and angles identical.) Each vertex of the solid is also identical to any other,i.e. has the same number of faces adjoining.

Show there can be no more than five Platonic solids. (There are in fact exactly five: the tetrahedron, octahedron, cube, icosahedron and dodecahedron.)

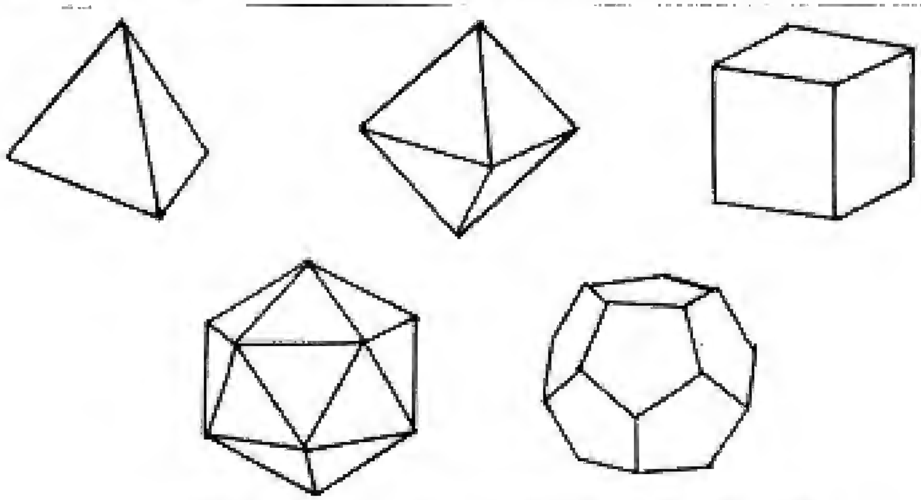

Fig. 31. Platonic Solids: Tetrahedron, Octohedron, Cube Icosahedron, Dodecahedron.

Hint: The surface of a Platonic solid is topologically a sphere so has total curvature 4π. This is distributed among v vertices (no curvature along any edge!) all containing the same amount of curvature, c.

$$vc=4\pi$$

Each vertex on the other hand is made up of a vertex from each of the f faces coming

together there. The vertex of a regular polygon has interior angle i.

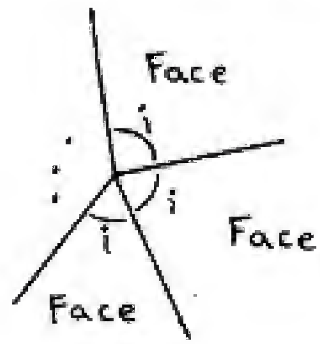

Fig. 32. A vertex of a solid where the vertices of the faces meet.

Thus

$$c = 2\pi - fi$$

(If you don't understand this, go over the section on curvature of cones.) There's one more formula I can write down. Each face has s sides and so

$$i = \pi - 2\pi/s.$$

Let's start with 3 sided faces. $i=\pi/3$, $c=2\pi-f\pi/3$. Now f can't be 1 or 2 because at least 3 faces must meet at each vertex (that's more or less obvious). c, however can't be zero or negative (since $vc = +4\pi$). So that only leaves f=3,4, or 5 with $v=4\pi/c=4/(2-f/3) = 12/(6-f) =$ 4, 6, or 12. These possibilities are tetrahedron, octahedron and icosahedron respectively.

Now how about 4 sided faces, squares. The only way to make a proper vertex out of squares is with 3 of them.

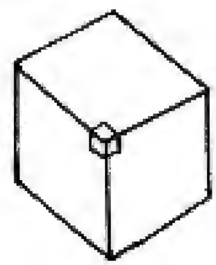

Fig. 33. Cube: $s=4$, $i=\pi/2$, $c=\pi/2$, $v=8$.

That's a cube.

Five sided faces: $s=5$, $i=\pi-2\pi/5=3\pi/5$. The only way to make a proper vertex out of such vertex angles is with 3 faces, $c=2\pi-9\pi/5=\pi/5$, $v=20$. That's a dodecahedron.

How about 6 or more sided figures? The basic problem is that 6 or more sided regular polygons have 120° or more at each vertex. And since you must have at least 3 of them at each vertex of the solid, you have at least 360° of "cone" at each of the solids vertices. That just won't work, because it makes negative curvature (see CONES again).

If you like arithmetic, try it this way: Remember $c>0$. But $c=2\pi-fi=2\pi-f(\pi-2\pi/s)$. This means

$$0<\pi(2-f(1-2/s))$$

and

$$0<2-f(1-2/s)$$

or

$$2/(1-2/s)>f.$$

But f must be at least 3 so

$$2s/(s-2)>3$$

$$2s>3s-6.$$

And finally we find that

$$s<8.$$

There are no Platonic solids with six or more sided figures as faces.

n+13:(A "new" excess) Suppose someone is unable to train his turtles to walk turtle lines any way except with their noses pointing in the direction of walk. He gives the following definition of the excess of a path.

$$\text{excess} = 2\pi - (\text{the turtle turning needed for the turtle to walk around the path})$$

He also says the turtle is turning positively when he's turning "toward the interior" of the path.

Can you make sense of that?

Is it the same excess as used in this paper?

Is it more precise or more useful or easier to understand than our standard definition?

To confuse the issue notice that the latitude 80° North path bounds <u>two</u> nice regions so it s excess should be the total curvature of both. Yet the regions obviously have different total curvatures. Can either or both definitions of excess explain this?

n+14:(Sewing) Does the dress of a dancer which is meant to have pleats even when she is spinning (and the dress is pulled outward by "centrifugal force") have positive or negative curvature?

An umbrella is made by sewing together six pieces shown roughly below.

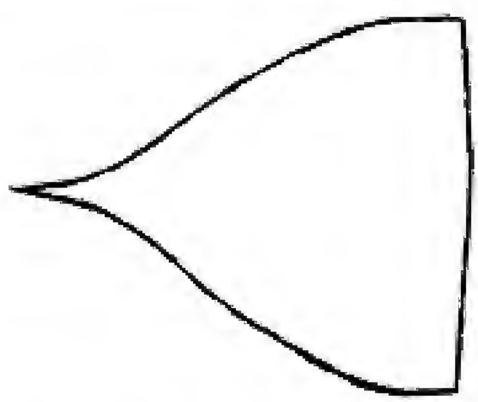

Fig. 34. A piece of an umbrella.

Where does the umbrella have negative curvature? How can you tell? Hint: Think about the finished umbrella. You can tell its total curvature from an edge ribbon (turtle road around the outside). If you start cutting away ribbon after ribbon from the edge, does the total curvature ever increase, indicating you cut away a strip of negative curvature? Now go back and relate this to things you could measure on the piece in Figure 34.

When you get done with this problem you should be able to look at Figure 34 and say, "Well, looks like the total curvature is almost exactly 2π, but there is some negative curvature from here to here."

n+15:(Broken turtle) A turtle is a bit out of adjustment and takes slightly longer (by 2%) strides with his left legs than with his right. (The distance between turtle's right and left legs is the same as his right side stride.)

What is the radius of the circle such a turtle would walk on a plane if he does not "turtle turn?" (Use turtle stride as your unit of distance.)

Suppose this turtle walks 25 steps and finds he has returned to his beginning position and heading. What is the total curvature interior to this path, presuming the interior is topologically nice?

How about a 100 step trip as above?

AKNOWLEDGEMENTS

I would like to thank Seymour Papert for telling me about the fellow who made the preposterous statement that turtles are no good on spheres. Thanks also to Hal Abelson for pointing out true, false, sloppy and crisp mathematics to me (thus guiding the development of this paper), and with whom the philosophy of elegant sloppiness which supported this effort was "cooked."